APPENDIX B

A Deeper Look at Moving Averages and the MACD

The charts and the discussion that follow may look a little strange, but the thought process is important. Too many times, traders use indicators or tools on their charts that they do not understand. Rules are developed based on the action of a squiggly line without fully understanding what that line measures and represents. Though some traders will find some success with an approach like this, it misses much of the potential in these tools. To fully understand a technical tool, it is useful to peer deeply into its construction, and to understand how it reacts to changes in the underlying market. One of the best ways to do this is to "feed" an indicator an artificially created dataset that focuses on specific types of market action. Think of this as a way to look at the indicator in a laboratory environment, isolating and controlling for various influences.

MOVING AVERAGES

If you want to really understand the tools you are using and how they will react to extreme situations, it helps to understand how they will react in the most simple, basic contexts as well. Of course, everyone grasps the basics of moving averages: add higher prices and the average will go up. When price flattens out, the average eventually will flatten, too. When price turns down, it will cut through the moving average, and, at some point, it will start pulling the average down with it. This is all simple, but it is not the point. The point is to build an intuitive and intimate understanding of the behavior of the average as it responds to changing market data. One of the best ways I have found to foster this intuition with any technical tool is to generate simple price patterns in artificial price series, and then to plot the indicator on this synthetic price data. (Nearly all charting packages allow users to import ASCII data, but if this is not possible, the indicator may be calculated directly in an Excel spreadsheet and graphed there.) Think of this as a

controlled laboratory experiment: you are controlling the data that is fed to the indicator so you can begin to understand the details of the indicator's reactions to that data. This process, along with careful thought and reflection, will build an understanding that goes beyond a simple understanding of the patterns on the screen.

The charts that follow are examples of this process. There is value in these specific examples, but it is even more important that you can take the procedure and adapt it to your own use. These examples show both 20-period *simple moving averages (SMAs)* and *exponential moving averages (EMAs)* to compare their behavior and fluctuations. The formulas required to build these indicators are already available from many books or the Internet, so they will not be a focus in this chapter, but you must have a clear conceptual understanding of the behavioral differences between the two. The simple moving average simply averages the price over a look-back window. It is completely blind to any data outside of that window, which creates the first potential issue: a simple moving average moves *twice* in reaction to any single large event. The value of the moving average has a change both when the event occurs and when it passes out of the left side of the evaluation window.

The EMA is a considerably more complex animal. Recent data are weighted more heavily in an EMA, and, technically, no data points are *ever* dropped from the average. Rather than being dropped, past data is rolled off with an exponential decay. In actual practice, the effect of distant past data far out of the evaluation period of the average is so small that it is insignificant, but it is still there. It is important to realize, though, that this effect smoothes the left-hand side of the evaluation window—an EMA will not jump twice as a simple moving average will. This is one of the main advantages of the EMA over the SMA.

Comparing EMA and SMA Behavior

Consider first how the moving averages react to a sudden shock in the market. Figure B.1 shows an artificial data series that is flat, then suddenly breaks into a precise linear trend, which just as abruptly comes to an end as prices flatten out again. Both the SMA (dotted line) and the EMA (solid line) are 20-period averages, but the front-weighting effect of the EMA causes it to react more quickly to the initial price shock. After a period of time, both averages settle into a steady relationship to the price trend, but the simple moving average is much quicker to return to the center after the trend stops. This effect is due to the decay in the EMA, which sees *all* data to the left of the average; the SMA is just a simple average of the past 20 data points. Once the market has been flat for 20 bars, the simple moving average exactly equals the close. There are two important lessons here. Most traders know the first, but few know the second. For an SMA and an EMA of the same length:

- The EMA will react faster to a large price change because it front-weights the data.
- However, the EMA is also slower to react to stabilizing prices because it has a very long look-back window.

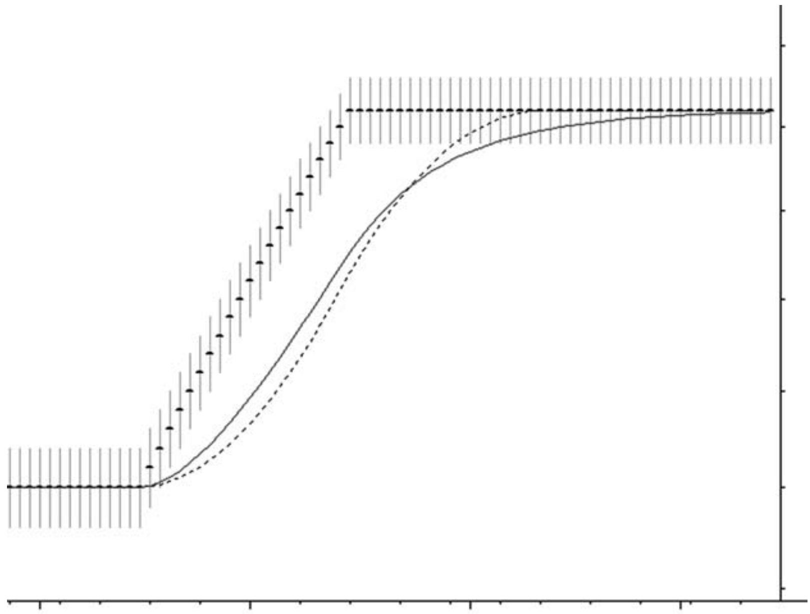

FIGURE B.1 Twenty-Period Simple (Dotted Line) and Exponential Moving Averages in a Market That Cleanly Shifts from Trading Range to Trend and Back

Many traders are convinced that moving averages provide support and resistance. Chapter 16 investigates this claim in some depth, and shows that there is no proof that any moving average is any better than any random number. One reason that so many traders remain convinced of this support/resistance effect is because it looks so convincing on charts. It is easy to find example after example of places where price touched a moving average and then shot away, but this is a result of two effects: One, we tend to attach more significance to lines on charts and to perceive patterns in random relationships. Two, there is a mathematical reason for this, as moving averages will approach prices as they pause in trends due to simple math. Figure B.2 illustrates this effect in an idealized market moving in stair-step trend legs.

If the rate of a trend is constant (arithmetically, not geometrically), a moving average will eventually settle into a consistent visual relationship with that market, tracking it at the same rate below, for an uptrend, or above in the case of a downtrend. If the rate of the trend increases, or if there is a shock in the opposite direction, the EMA will react before the SMA, but, as before, the SMA will react more quickly to stabilizing prices. It is also not well known, but the EMA will approach price more closely than a SMA in a stable trend, again due to the front-weighting in the EMA calculation. Essentially, the EMA catches turns a bit faster, but is much slower to come into the new, stable value when the trend ends and transitions into a new trading range. This is expected behavior, as the EMA should respond quickly to recent data while maintaining a memory of the long data history. Figure B.3 shows a trend with two inflection points and a clear ending

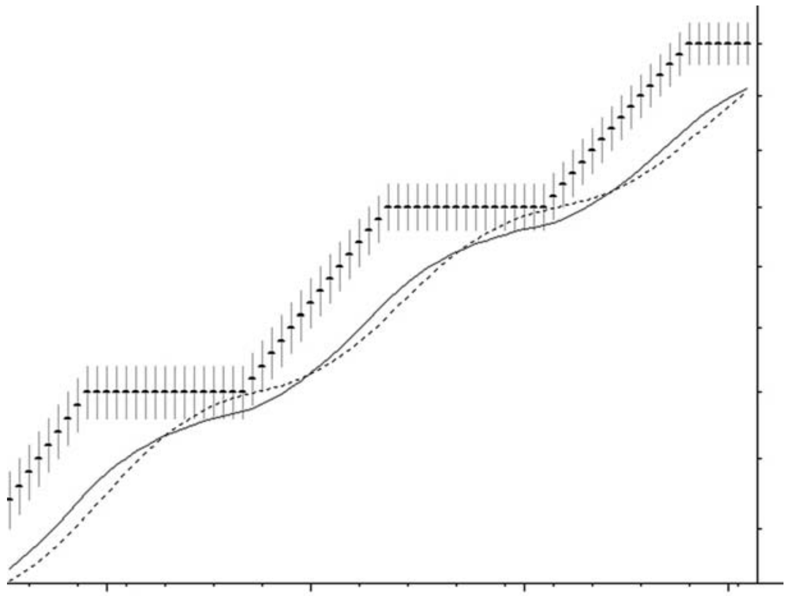

FIGURE B.2 Are the Moving Averages Supporting Prices?

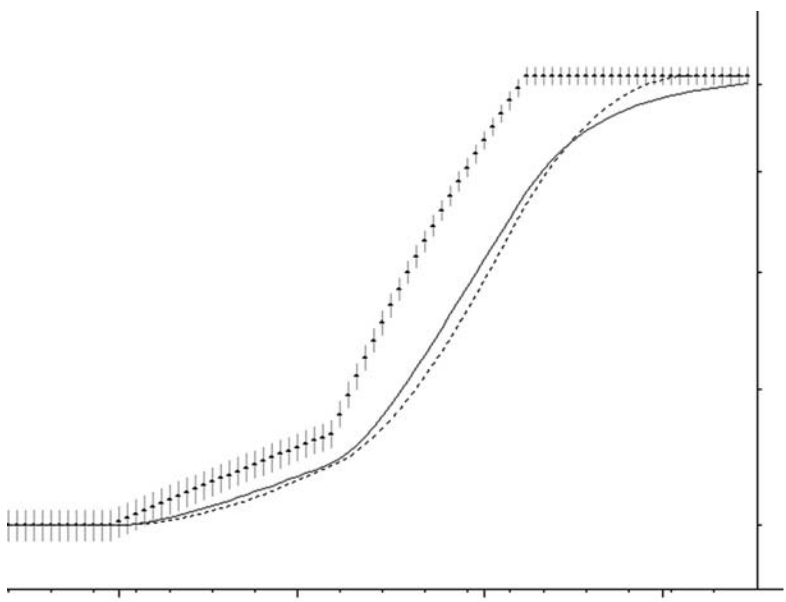

FIGURE B.3 Notice the Different Behavior of the EMA and the SMA at Inflection Points and at the End of the Trend

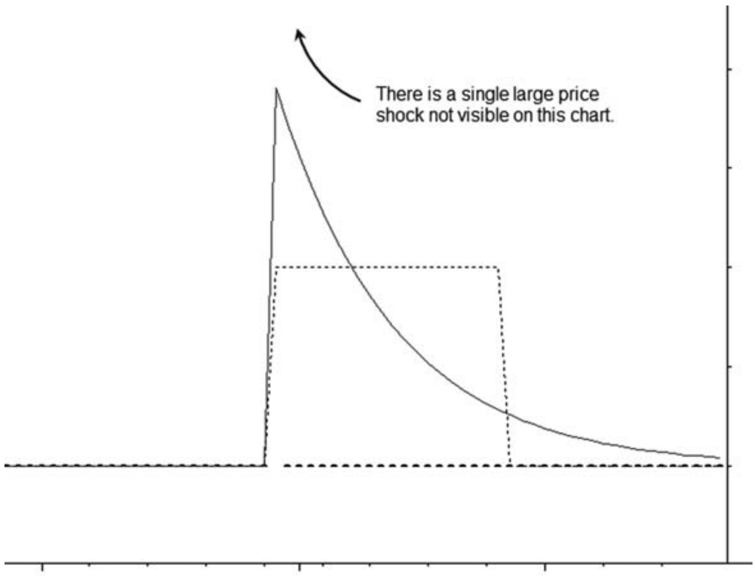

FIGURE B.4 Notice the Different Behavior of the EMA and SMA in Response to a Single Large Price Shock

point. Spend some time thinking about how the SMA and EMA react differently at these inflections; if you are using the slope of a moving average, or a moving average crossover, to define trends, you need to be aware of these issues. One average is not better than the other, but you need to be mindful of the differences.

Figure B.4, which has a single price bar that is not visible far above the top of the chart, is useful for building intuition about how the averages react to a sudden shock. Note that the SMA shows *two* inflection points when there was actually only *one* event on the chart. The second inflection (the drop) in the SMA was merely an artifact created as the price spike moved out of the average's evaluation window. In this particular case, the EMA probably more accurately reflects what is going on in the market. Traders using an SMA in an intuitive fashion are not likely to be misled, because they are focusing on the bigger picture, but systematic approaches or tools (such as trend indicators) derived from an SMA may have some issues with outliers. Particularly in intraday data, where the overnight gap is significant, or longer-term equities, which have frequent price shocks due to earnings announcements, systematic tools based on simple moving averages are subject to distortion.

Few traders realize that simple moving averages are low-pass filters, meaning that they will filter out (eliminate) higher-frequency oscillations and cycles. Figure B.5 shows a situation that will probably not be encountered in actual trading: a market that is moving in an idealized 20-period sine wave, with 20-period SMA and EMA applied. Though it may be very counterintuitive, the SMA is completely flat. When the SMA length matches the sine wavelength, there are always as many values above as below the moving average,

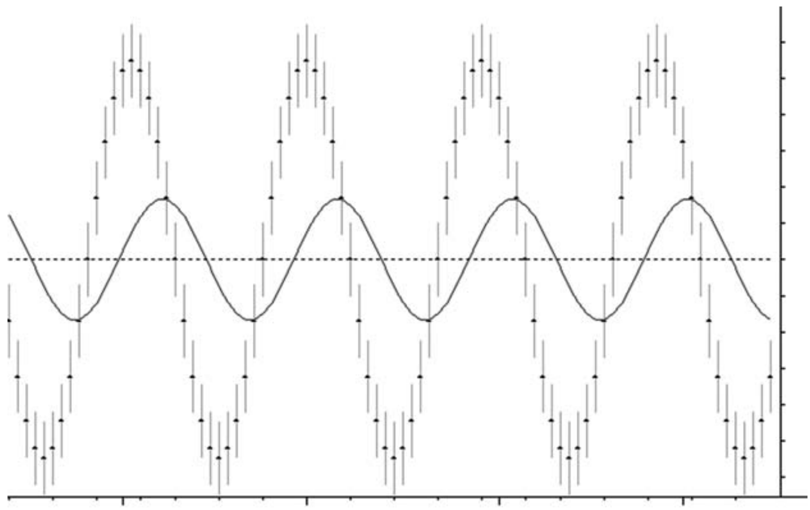

FIGURE B.5 Twenty-Period EMA and SMA Applied to a Theoretical 20-Period Sine Wave
Notice that the SMA (dotted line) is completely flat on the zero line.

so the average nets to zero—an SMA will always hide any cycles that are whole-number multiples of the SMA length. You might think that this is pure theory and that it would not have any application to a real trading situation, but extraordinary situations do occur. I once had a trader I was working with point out that his moving average on a 1-minute chart was not "working right" on part of the chart because it was not reacting to the market's movements. For about 45 minutes, this stock had settled into a very dependable cycle that just happened to match the length of the moving average he was using, so, while the market was oscillating fairly wildly, his average wouldn't budge. There are cycles in the market. It is difficult to trade them because they are ephemeral and they shift frequently, but they will sometimes line up with your moving averages with seemingly bizarre results. Trade long enough and you will see pretty much everything.

Let's end this section with an example that is important for longer-term investors to keep in mind. The natural language of the market is percentage changes and growth rates, which is why finance math is based on discounting cash flows and compound interest, and why the first task of any research project is converting prices into returns. Shorter-term traders tend to think in differences (e.g., "I made a point and a half in that stock."); longer-term investors think more often in percentages. Figure B.6 shows a market that is growing at a constant rate; each data point is a 5 percent increase over the previous one. On a linearly scaled chart, a market appreciating at a constant growth rate will describe a *curve*, not a straight line. On a linear chart, moving averages will seem to lag behind the price curve at an ever-increasing distance.

What is actually happening, however, is that the moving averages lag a *constant percentage* behind prices. The linear chart does not misrepresent anything, but it is not the right tool to look at percentage-based relationships. Figure B.7 is exactly the same as

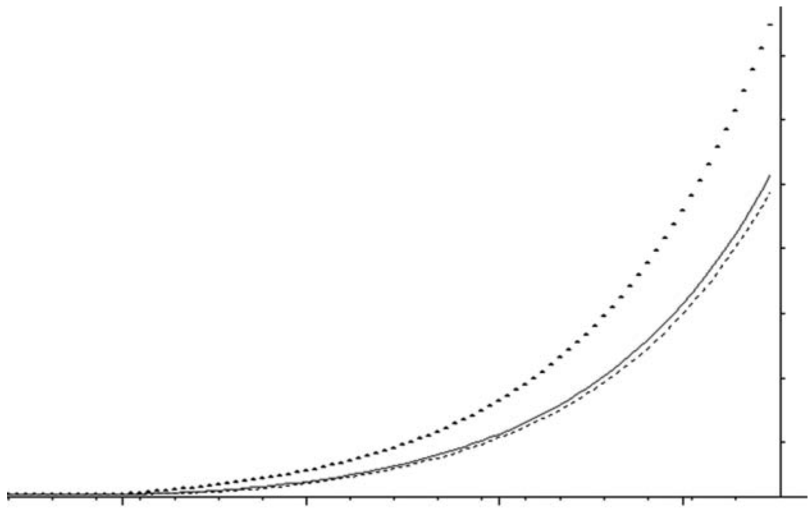

FIGURE B.6 Moving Averages Applied to a Market Growing at 5 Percent Constant Rate, Linear Scale

the previous chart; the only difference is that the y-axis is log scale rather than linear. Remember, log scale charts are designed so that equal distance intervals on the y-axis are equal percentage changes, not equal price changes as on a linear chart. Note the ticks on the right side of the y-axis, which are evenly spaced prices on the linear chart, become compressed near the top of the axis when log scaled. Be clear on this effect: straight lines on log charts are curves on linear charts. If you are drawing trend lines on

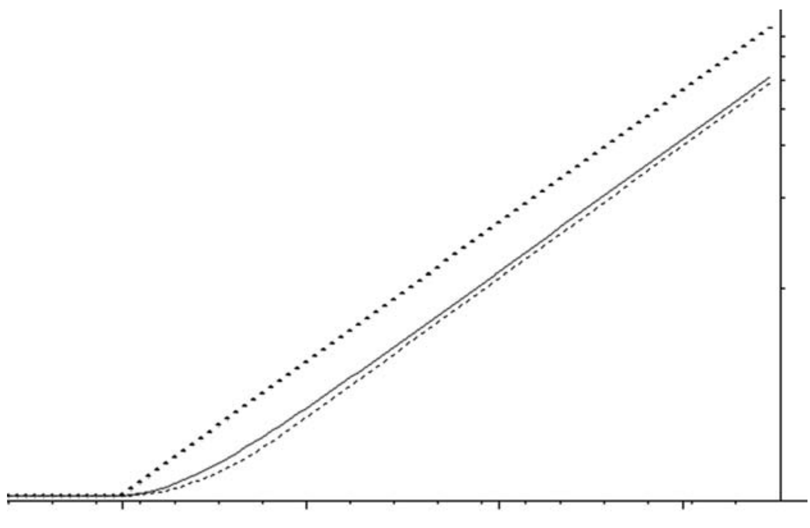

FIGURE B.7 Moving Averages Applied to a Market Growing at a Constant 5 Percent Rate, Log Scale

Note that curves on previous chart are straight lines here.

log scale charts, you are actually drawing curves on normal, linearly scaled charts. This may be perfectly correct in the context of a longer-term chart where the curve faithfully represents the growth rate, but it is something that must be fully understood.

THE MACD

Moving average convergence/divergence (MACD) was one of the early technical tools developed by Gerald Appel in the mid-1970s. As an interesting aside, this is one of the few commonly used indicators whose name actually says something accurate about the tool. Consider some other indicators: stochastics have nothing at all to do with stochastic processes; the Relative Strength Index (RSI) neither measures relative strength nor is it an index; the Commodity Channel Index (CCI) is not commodity-specific, nor does it deal with channels, nor is it an index. This is not an exhaustive list, but you get the idea. The standard MACD consists of four elements: a fast line (the MACD proper), a slow line (the signal line), a zero line for reference, and a histogram (bar chart) that shows the difference between the slow and fast lines. The standard MACD is constructed from exponential moving averages, but this modified version uses simple moving averages and dispenses with the MACD histogram altogether, resulting in a cleaner indicator.

Basic Construction of the MACD

The fast line of the MACD measures the distance between a shorter-term and a longer-term moving average. To understand exactly what this measure says about price action, think about how different periods of moving averages will respond to price movement. A moving average with a short look-back window (period) will track price movements more closely than an average with a longer period. Figure B.8 shows 3-period and 10-period simple moving averages applied to a daily chart of two-year Treasury notes. The line plotted below the price bars is the fast line of the MACD, which is simply the value of the slow moving average subtracted from the value of the fast average. When the fast average is above the slow average, this line is positive and vice versa. Notice the important behavior at the points marked A and B on this chart. Though price was higher at B, the distance between the moving averages was actually smaller, so the indicator registers a lower level at B. The distance between these two averages is one way to measure the momentum behind a market's movements, and this lower peak in the indicator suggests that the second price high was made on lower momentum.

Another important point is that the fast line of the MACD will register zero when the two averages cross; this highlights a condition of relative equilibrium on the time frame being measured. Figure B.9 marks spots where the fast line crosses the zero line, and shows that this happens when the moving averages intersect.

Figure B.10 shows the slow line (sometimes called the signal line) of the MACD, which is simply a 16-period moving average of the fast line. It is important to note that,

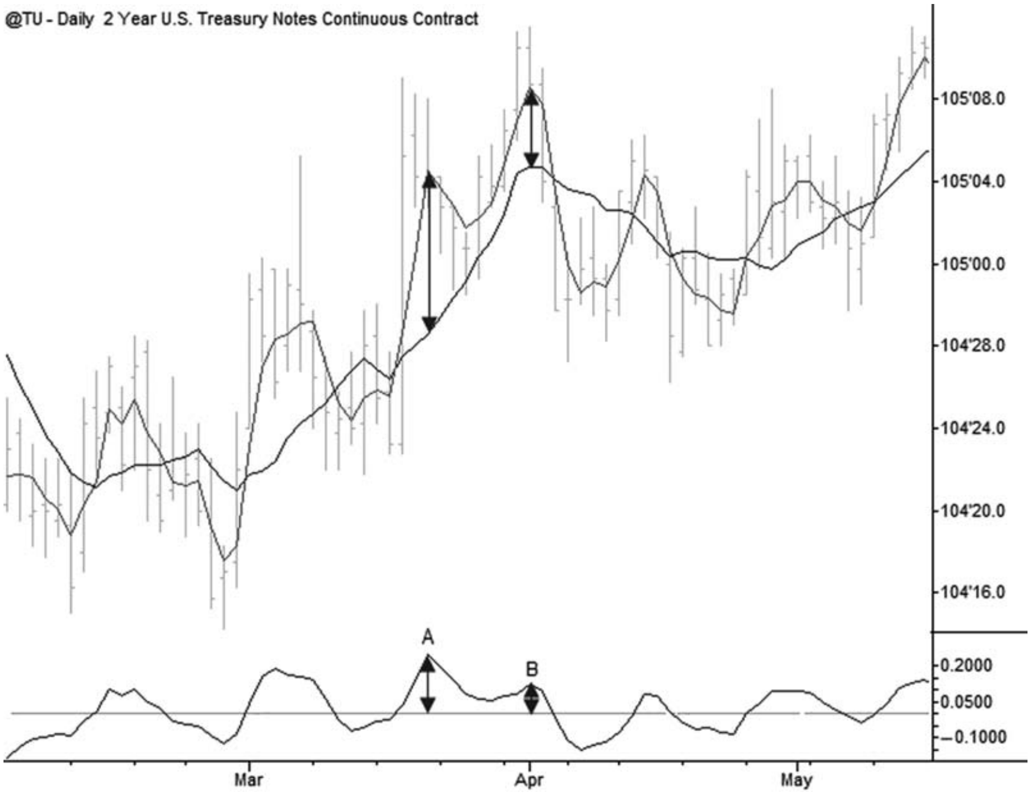

FIGURE B.8 The Fast Line of the MACD Measures the Distance Between Two Moving Averages

being a smoothed version of the fast line, it lags the fast line considerably, but also generally reflects the trend of the fast line—if the fast line is far above zero, the slow line will usually be sloping upward and vice versa. There are several ways to use the slow line, but the general concept tying them all together is that it reflects the trend on an intermediate-term time frame.

This is the construction of the particular variation of the MACD that I use. For each bar, calculate:

- Fast line: 3-period SMA minus 10-period SMA.
- Signal line: 16-period SMA of the fast line.
- Histogram: none.
- Plot a zero line for reference.

Using the same (3, 10, 16) settings in a standard MACD will give similar results, but, in my experience, the long memory of the exponential moving averages (EMAs) does make a difference at times. You will certainly be able to apply the same concepts to the standard MACD, but make sure you understand the differences between the two

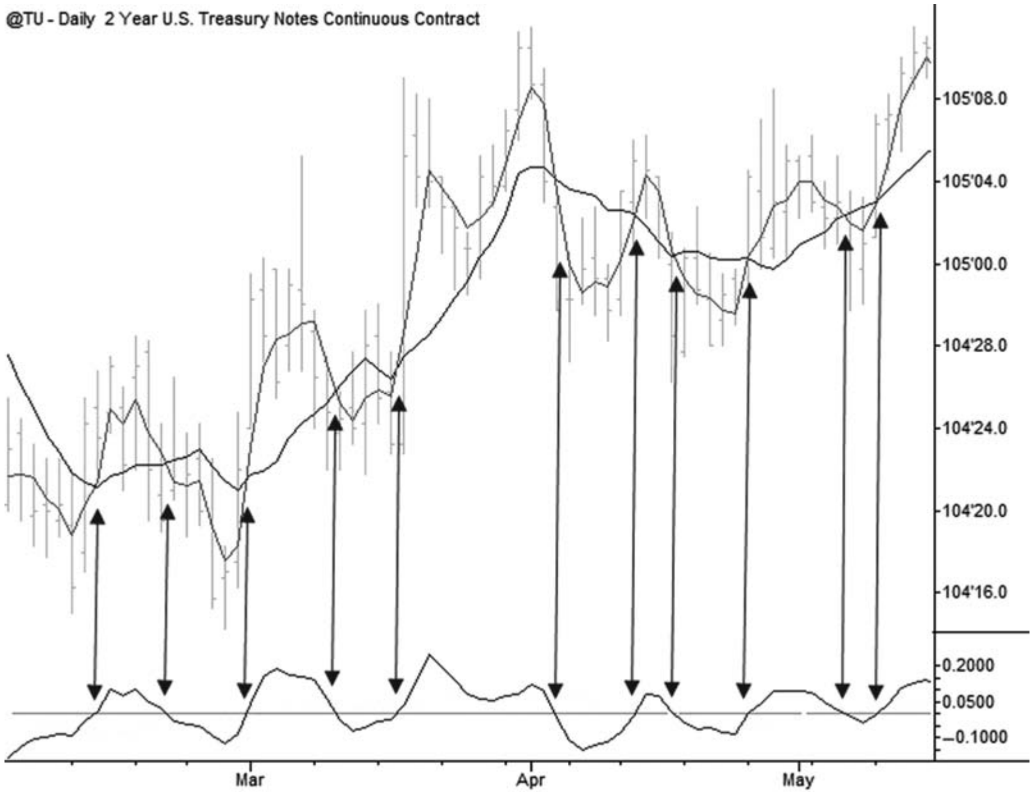

FIGURE B.9 The Fast Line of the MACD Crosses Zero When the Moving Averages Touch

indicators. Figure B.11 compares the modified with the standard MACD (3, 10, 16). Notice that the standard tool adds another plot, usually plotted as bars behind the indicator. This plot, the MACD histogram, shows the difference between the fast and slow lines of the MACD histogram, so it is actually the MACD of the MACD. Some traders find this to be a useful component, but it is possible to extract much of the same information from a careful reading of the fast and slow lines themselves.

A Deeper Look

The fast line is very sensitive to changes in the rate of change of prices. Read that again, carefully: the fast line swings up in response to the second derivative, or the rate of change of the rate of change of price. When we actually work with this tool, we usually think of it a little more loosely, as simply measuring the momentum of prices, but it is a good idea to be as precise as possible here at the beginning—this tool measures *changes* in momentum, not momentum itself. To begin to build some intuition about this tool, Figure B.12 shows a modified MACD applied to an idealized price series that breaks into

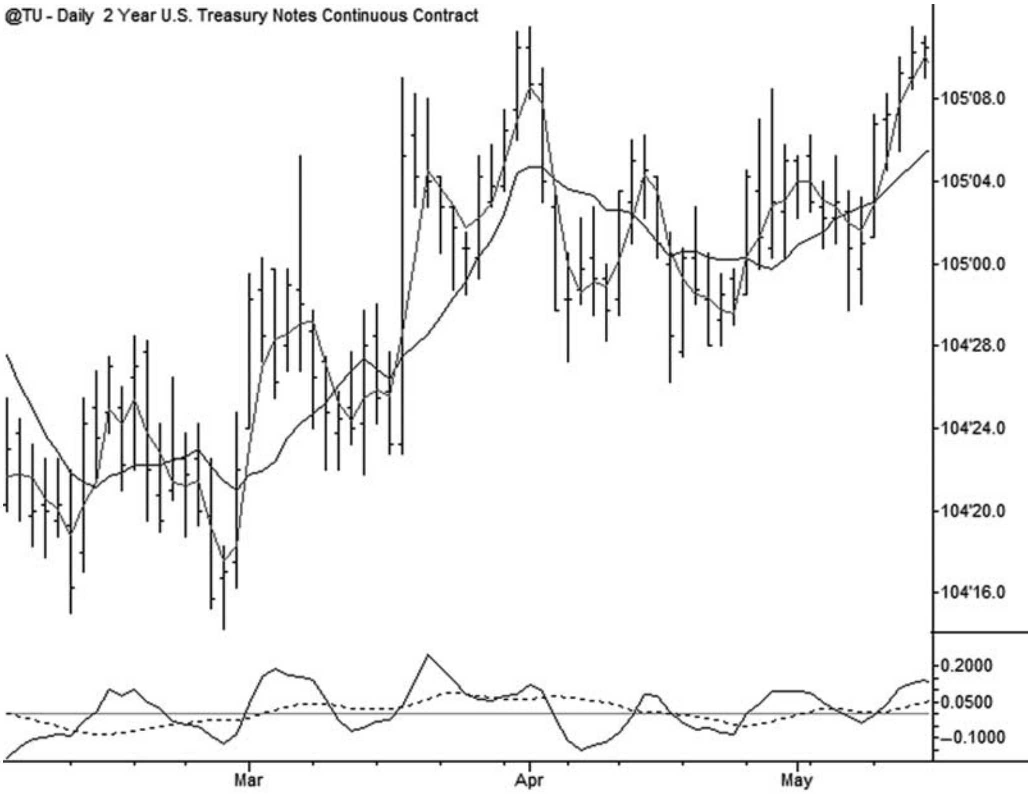

FIGURE B.10 The Slow Line of the MACD Is a Smoothed Version of the Fast Line

a steady-state uptrend and then cleanly transitions to a sideways market. Consider the inflections in the MACD:

- A: The fast line responds immediately to the change in the market by hooking higher on the first price bar of the uptrend.
- B: Everything between points A and B is an artifact of the indicator. Though the slope of the fast line changes in a curve, this does not reflect any change in the rate of trend in the market, which is trending steadily higher. At B, 10 bars into the uptrend, the 10-period moving average is now trending steadily with prices (remember, the fast line measures the difference between a 3-period and a 10-period moving average), so the indicator goes flat.
- C: The fast line again responds immediately to the shift in the market by hooking down on the first bar that breaks the trend pattern. Note that the MACD fast line going down does not mean that prices are going down, but that the rate of change of prices has gone down, in this case to zero.
- D: The fast line levels out, again 10 bars following the change in the market.

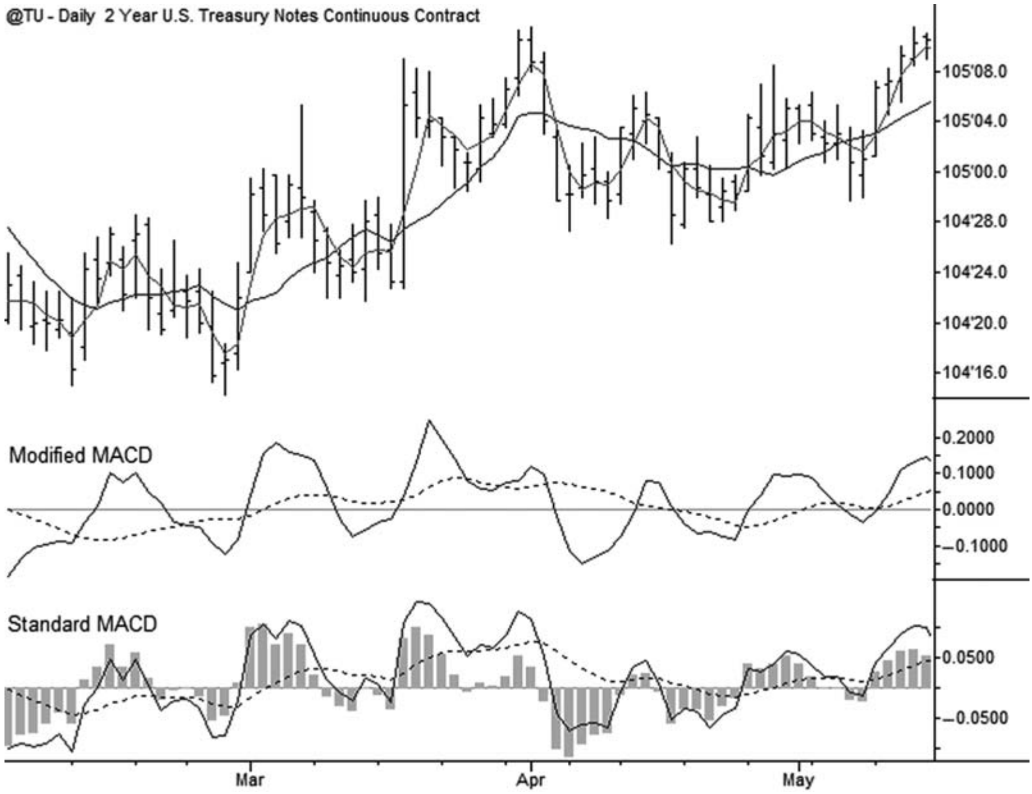

FIGURE B.11 Modified and Standard MACDs (Both 3, 10, 16) Compared

If you are looking at Figure B.12 and thinking that this could never be relevant in actual trading, consider the daily chart of Apple Inc. in Figure B.13. After a downtrend, the market transitioned to an uptrend at the point marked A, and the MACD fast line immediately responded by hooking upward, just as in the inflection marked A in the previous chart. Real market action, of course, contains much more noise and variation than our simplified example, but much of the downturn in the line at point B in Figure B.13 is due to this 10-bar artifact. A trader paying too much attention to the indicator's line at this point might surmise that momentum had turned downward. Prices did take a small pause at this point, but the indicator's reaction was out of proportion to that change. Point C is also interesting; a large move up is required to hook the fast line up after an extended trend. Half of the battle with using indicators is knowing when to use them and when to ignore them. If you are reacting to every jot of the indicator, you are missing the point. It is much better to use the tool only at potential inflections to add another layer of confidence to analysis that focuses primarily on the price bars themselves.

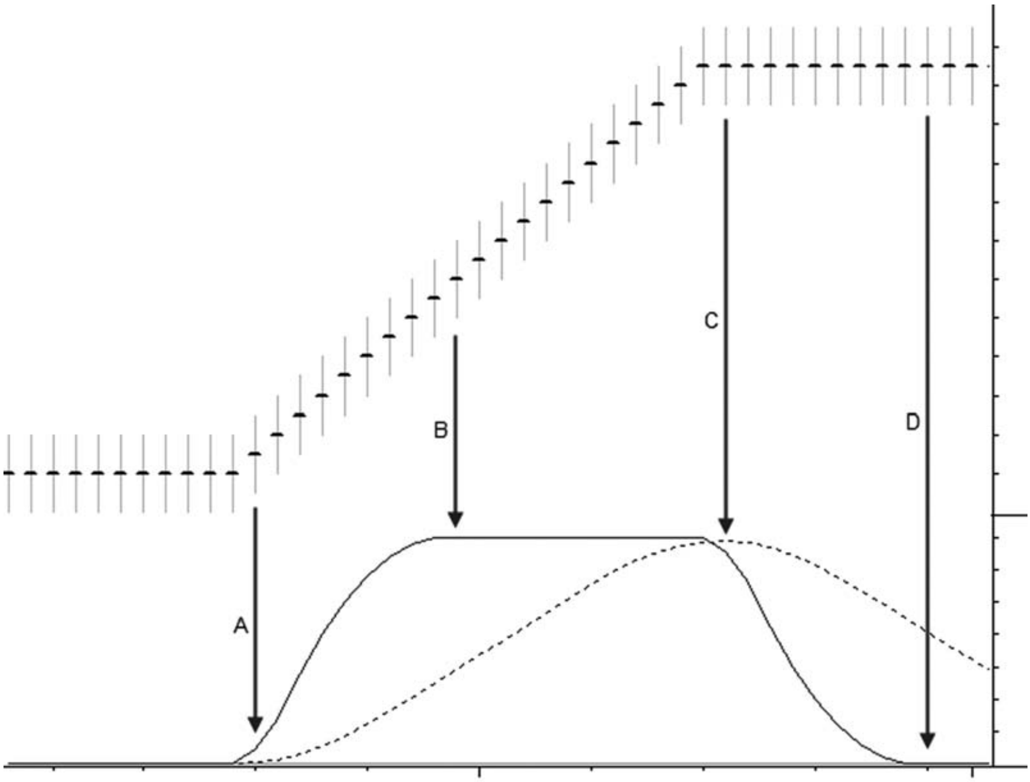

FIGURE B.12 Note That Inflections in the Fast Line of the MACD Include Some Artifacts Due to the Indicator's Construction

The modified MACD is constructed from simple moving averages, so it suffers from the problem of dual inflections to single price shocks. In fact, the multiple moving averages result in a complex reaction to single large events. This is rarely a significant problem when we are actually using the tool, but it can be an issue in some situations—for instance, in applying the tool to intraday data following a large overnight gap. Figure B.14 shows a situation that would be absurd in a real market: an otherwise absolutely flat market with a single large price shock, immediately reversed the following day. Note that the MACD's two lines have 12 inflections in response to this single bar! Again, this is not a tradable feature, but it does highlight the folly of trying to follow every move of the indicator too closely. In a real market situation, the indicator is nearly always irrelevant after a large price shock, and attention should be focused on other factors such as price action following the movement.

So far, these theoretical data sets have been so clean that it may be difficult to relate them to situations that are likely to occur in actual trading. Consider next what happens when the MACD is applied to a data set that begins to more closely approximate real market conditions, as in Figure B.15. This data set alternates uptrending and

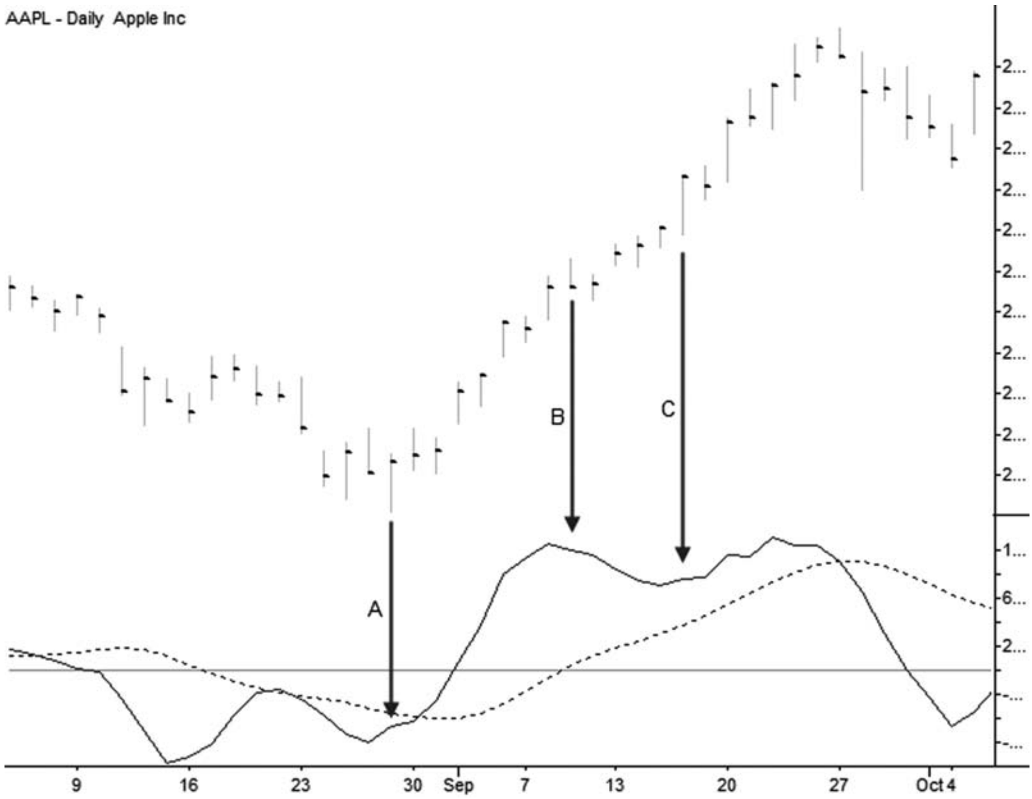

FIGURE B.13 Consider the Inflections of the MACD Applied to an Actual Market

downtrending legs, though it reduces them to consistent, idealized linear changes. Specifically, the down legs (BC and DA) are the same rate and length, while the up leg CD moves at a rate equal to 75 percent of AB's changes. The indicator registers a lower peak at D compared to B, which would suggest to many traders that trend leg CD, though it made a new high relative to the point B, did so on lower momentum; this is an example of a so-called *momentum divergence* on the MACD. This example shows precisely what lower momentum means for the MACD: here it means that the lower-momentum leg had smaller daily changes compared to the higher-momentum leg. Note that, in this case, both legs were the same length in terms of number of bars, so the lower-momentum leg, overall, covered a smaller range of prices than the higher-momentum leg. If the lower-momentum leg had continued for more days it would eventually have moved the same distance as the higher-momentum leg, but it still would have shown a divergence with the indicator at the second peak.

There is, however, another possibility, shown in Figure B.16, which has two uptrending legs (AB and CD) that both move with the same rate of change. Here, the second leg (CD) is shorter, including fewer bars than the longer leg, and the indicator again registers

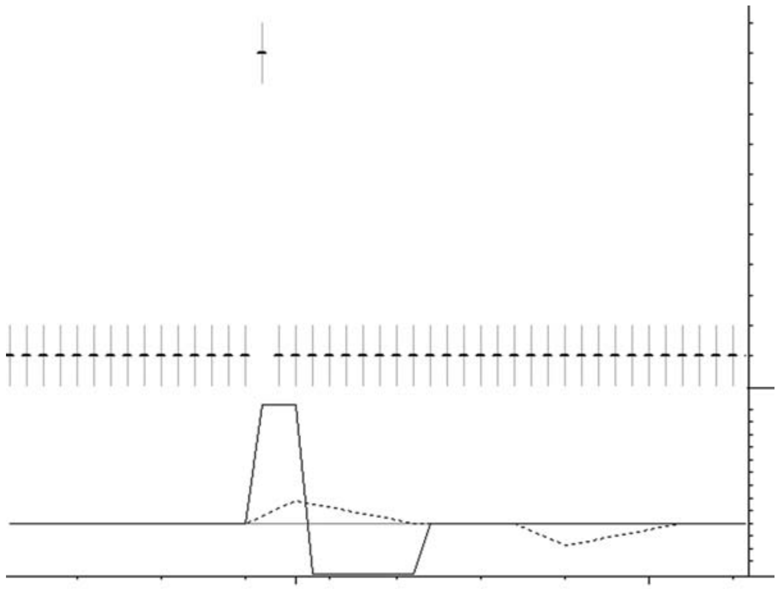

FIGURE B.14 Multiple Inflections on the MACD Following a Single Price Shock

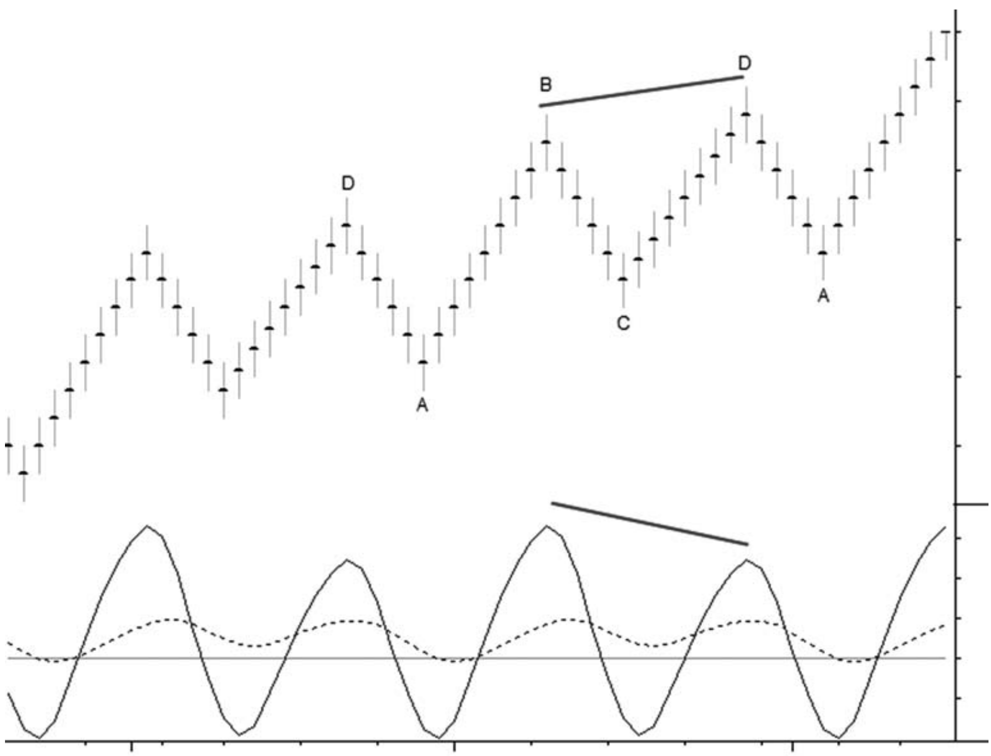

FIGURE B.15 Momentum Divergence Due to Lower Rate of Change on the Second Leg

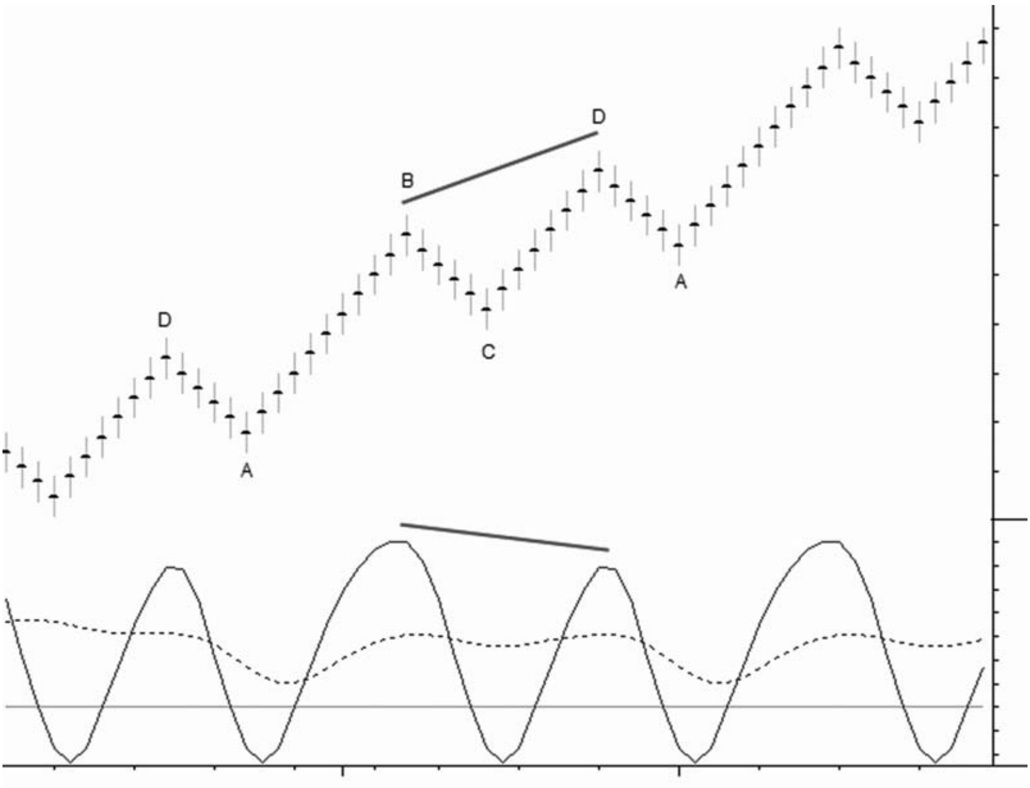

FIGURE B.16 Momentum Divergence Due to a Shorter Second Leg

a divergence. This is fundamental to divergence on momentum indicators: divergences will register with trend legs that either move at a lower rate of change or extend for fewer bars relative to the higher-momentum leg. In actual practice, the presence of additional noise and fluctuation obscures these simple tendencies, but this is the underlying truth of momentum divergence. Very few traders think about these concepts with precision.